Experimental and causal view on information integration in autonomous agents

Philipp Geiger[1] and **Katja Hofmann**[2] and **Bernhard Schölkopf**[3]

arXiv:1606.04250v2 [cs.AI] 26 Aug 2016

Abstract. The amount of digitally available but heterogeneous information about the world is remarkable, and new technologies such as self-driving cars, smart homes, or the internet of things may further increase it. In this paper we examine certain aspects of the problem of how such heterogeneous information can be harnessed by autonomous agents. After discussing potentials and limitations of some existing approaches, we investigate how *experiments* can help to obtain a better understanding of the problem. Specifically, we present a simple agent that integrates video data from a different agent, and implement and evaluate a version of it on the novel experimentation platform *Malmo*. The focus of a second investigation is on how information about the hardware of different agents, the agents' sensory data, and *causal* information can be utilized for knowledge transfer between agents and subsequently more data-efficient decision making. Finally, we discuss potential future steps w.r.t. theory and experimentation, and formulate open questions.

1 Introduction

Increasing amounts of heterogeneous information are recorded and connected, and this trend is likely to continue in the light of new technology such as self-driving cars, smart homes with domestic robots, or the internet of things. Intuitively, it makes sense to design autonomous agents in a way that they automatically integrate all relevant and well-structured information on their environment that is available. Various aspects of the problem of designing such agents have been investigated previously. In this paper we approach the problem from two directions which, to our knowledge, have not been (exhaustively) examined yet: using sophisticated simulated experiments, on a practical level, and causal models, on a more theoretical level. The complexity of the problem allows us only to take small steps.

1.1 Main contributions

The main contribution of this paper consists of two investigations:

- In Section 5 we use a simulated *experimentation platform Malmo* to obtain a *better understanding* of the problem of integrating heterogeneous information. More specifically, we present a simple agent that harnesses video data from a different agent, and implement and evaluate a version of it.
- In Section 6 we investigate how detailed information on the *hardware of different agents* (we consider self-driving cars as example), their sensory data, and physical or *causal information* can be utilized for knowledge transfer between them and subsequent more data-efficient decision making.

The common structure of both investigations is that we start with a description of a scenario that captures certain core aspects of the general problem, in particular containing a variety of heterogeneous information sources, and then sketch a method to perform information integration and subsequent decision making in these scenarios. After experimentally evaluating the method, or illustrating it based on a toy example, we conclude both investigations with a discussion of the advantages and limitations of the respective methods.

A reoccurring theme in our investigations is that we try to treat as much information (including models) as possible explicitly as input to algorithms instead of implicitly encoding it into algorithms. Our hope is that this sheds a better, more explicit light on the problem.

1.2 Structure of the paper

The paper is structured as follows: We introduce the experimentation platform and basic concepts in Section 2. In Section 3, we formulate the problem. In Section 4, investigate potentials and limitations of existing approaches for the problem. In Sections 5 and 6, we present our two main investigations. In Section 7, we discuss future directions and pose open questions. We conclude with Section 8.

2 Preliminaries

Here we introduce the concepts, models and the experimentation platform we will use in the paper.

Autonomous agents. By an *(autonomous) agent* we mean a mechanism which, at each time point t, takes some input from the *environment*, in particular its sensory data which we refer to as *observation*, and outputs some *action* that influences the environment. Moreover, by an *intelligent (autonomous) agent* we mean an autonomous agent which is successful in using its inputs and outputs for given *tasks*, i.e., specific goals w.r.t. the environment, often encoded by a *reward* or *utility* function.

Note that in this paper we do not define a clear boundary between agent and environment. Usually, we consider the hardware platform of an agent (e.g., the car) as part of the agent. This particularly has to be kept in mind when we talk of several agents in the "same environment": the hardware of the agents may still differ. (It is almost a philosophical problem to define what precisely the "same environment" means. Here we simply suggest to interpret this notion as if it

[1] Max Planck Institute for Intelligent Systems, Tübingen, Germany; email: philipp.geiger@tuebingen.mpg.de
[2] Microsoft Research Cambridge, Cambridge, United Kingdom; email: katja.hofmann@microsoft.com
[3] Max Planck Institute for Intelligent Systems, Tübingen, Germany; email: bs@tuebingen.mpg.de

were used in an everyday conversation. We pose a related question in Section 7.)

When we consider some agent A w.r.t. which some task is given and for which we want to infer good actions, while information may come from (sensory data of) a collection C of other agents, then we refer to A as *target agent* and the agents in C as *source agents*.

Experimentation platform "Malmo". For the experiments in Section 5 we will use the software *Malmo*, a simulated environment for experimentation with intelligent agents, that was introduced recently [4]. Malmo is based on "Minecraft", which is an open-ended computer game where players can explore, construct, collaborate, and invent their own "games within the game" or tasks. The Malmo platform provides an abstraction layer on top of the game through which one or more agents observe the current state of the world (observations are customizable) and interact with it through their specific action sets (or actuators). The advantage of Malmo is that it reflects important characteristics of the problem instances we will introduce in Section 3. To illustrate the platform, three sample observations of an agent in three different maps in Malmo will be depicted in Figure 1.

Causal models. Mathematically, a *causal model* [10, 14] M over a set V of variables consists of a directed acyclic graph (DAG) G with V as node set, called *causal diagram* or *causal DAG*, and a conditional probability density $p_{X|PA_X=pa_X}$ (for all pa_X in the domain of PA_X) for each $X \in V$, where PA_X are the parents of X in G. Given a causal model M and a tuple of variables Z of M, the *post-interventional causal model* $M_{\mathrm{do}\ Z=z}$ is defined as follows: drop the variables in Z and all incoming arrows from the causal diagram, and fix the value of variables in Z to the corresponding entry of z in all remaining conditional densities. Based on this, we define the *post-interventional density of* Y *after setting* Z *to* z, denoted by $p_{Y|\mathrm{do}\ Z=z}$ or $p_{Y|\mathrm{do}\ z}$, by the the density of Y in $M_{\mathrm{do}\ Z=z}$.

On a non-mathematical level, we consider M to be a correct causal model of some part of reality, if it correctly predicts the outcomes of *interventions* in that part of reality (clearly there are other reasonable definitions of causation). Keep in mind that in this paper, in particular Section 6, will use causal models and causal reasoning in a more intuitive and sometimes less rigorous way, to not be limited by the expressive power of the current formal modeling language.

Note that we will use expressions like $p(x|y)$ as shorthand for $p_{X|Y}(x|y)$.

3 Problem formulation

Let us describe the problem we consider in this paper in more detail:

- *Given:* a task T w.r.t. some partially unknown environment E, and additional heterogeneous but well-structured information sources H (e.g., in the form of low-level sensory data, or in the form of high-level descriptions).
- *Goal:* design an agent A that automatically harnesses as much relevant information of H as possible to solve T; more specifically, it should use H to either improve an explicit model of the effects[4] of its actions, which then guides its actions, or let its actions directly be guided by H.

[4] In this sense, at least the target of the information integration is clear: modeling the dynamics or causal structure of the agent in the environment.

Note that alternatively, one could also formulate the problem by letting A only be an actuator, and not a complete agent, and include the agent's sensors into H. This might be a more precise formulation, however, for the sake of an intuitive terminology, we stick to the definition based on A being an agent.[5]

To illustrate the general problem, in what follows, we give three concrete examples of desirable scenarios in which agents automatically integrate heterogeneous information. Ideally, agents would be able to *simultaneously integrate information sources from all three examples*.

3.1 Example 1: sharing information between different self-driving cars

Consider self-driving cars. It is desirable that as much information about the environment can be shared amongst them. By such information we mean up-to-date detailed street maps, traffic information, information on how to avoid accidents etc. For instance, assume that one self-driving car leaves the road at some difficult spot due to some inappropriate action, since, for instance, the spot has not been visited by self-driving cars before (or newly appeared due to say some oil spill or rockfall). If we only consider other self-driving cars of the same hardware, this experience could directly be transferred to them by enforcing them not to perform the very action at the very spot. (I.e., for all cars of the same hardware one could treat the experience as if it was their own and make them "learn" from it in the usual reinforcement learning (RL) way.) However, if we assume that there are self-driving cars of different types, then it is not possible to transfer the experience, and thus avoid further accidents, in this straight-forward way.

3.2 Example 2: observing another agent

Consider domestic robots. A domestic robot may, with its sensor, observe humans how they handle doors, windows, light switches, or kitchen devices. It should be possible that domestic robots learn from such experience. For instance, one could imagine a robot to reason that, if it is able to operate the door knob in a similar way as a human did before, this would also open the door and and thus allow the robot to walk into the other room (to achieve some task).

3.3 Example 3: integrating high-level information

Consider an agent that arrives in a city it has never been to before. The goal is to get to a certain destination, say to the town hall. A resident may be able to explain the way in a simple language, with words such as "... follow this street until you come to a church, then turn right ...". Or a resident could provide a map and mark directions on the map. One could imagine that an autonomous agent could combine such a description with a model of the "local" (or "low-level") dynamics that is shared by most environments (which is closely related to the laws of physics). The model of the "local" dynamics could have been either hard-coded, or inferred based on exploration in other (related) environments. In principle it should be possible that

[5] Note that the problem we formulate here does not coincide with developing ("strong") artificial intelligence (AI), as defined, e.g., by the Turing test or simply based on human-level intelligence. We restrict to sources of information that are more or less well-structured - either quantitative measurements with a simple and clear relation to the physical world, or information in a language much more restrictive than natural language. Nonetheless, the formulated problem can be seen as one step from say RL into the direction of AI.

such a combination allows the agent to successfully navigate to the destination city hall.

4 Related work: potentials and limitations

Various research directions exist that address major or minor aspects of the problem formulated in Section 3. Here we discuss the most relevant such directions we are aware of, highlighting their potentials and limitations w.r.t. the problem. Keep in mind that Sections 5.4 and 6.4 contain additional discussions on the advantages of our approaches over these directions.

4.1 Reinforcement learning

One of the most powerful approaches to shaping intelligent autonomous agents is *reinforcement learning (RL)* [16]. Instead of explicitly hard-coding each detail of an agent, for each environment and objective individually, the idea is to take an approach which is more modular and based on learning instead of hard-coding: the supervisor only determines the reward function and then the agent ideally uses exploration of the unknown environment and exploitation of the gained experience (sensory data) to achieve a high cumulative reward.

Regarding the problem we consider in this paper, RL plays a key role for integration of information in the form of *recordings of an agent's own past*, or of an agent with the same hardware. However, as mentioned in Section 3, in contrast to RL, here we consider the problem of integrating information beyond such recordings, such as sensory data from agents with different hardware, or higher level information such as maps.

4.2 Learning from demonstrations

According to [1], in *learning from demonstrations (LfD)*, some "teacher" performs a trajectory which is recorded, and the goal is that a "learner" agent, based on this recording, infers and imitates (or utilizes) the teacher's "policy" (or the dynamics of the environment, or both). Central notions that [1] uses to analyze and distinguish various types of LfD problems are the *record mapping*, i.e., what aspects of the teacher's demonstration are measured and recorded, and the *embodiment mapping*, i.e., if the recorded actions can directly be implemented by the "learner" and lead to similar observations as the recoded ones, or if the recordings first have to be transformed "to make sense" for the learner.

Our problem formulation can be seen as a generalization of LfD. Based on this, while a significant part of the problem we consider can be addressed by LfD methods, others are beyond the scope of these methods: Instead of hand-crafting, e.g., the embodiment mapping for each agent individually, we aim at (semi-)automating the inference of the mapping from recordings of "source" agents to actions of a "target" agent. In particular, we propose to do such a (semi-)automation based on additional information sources on the hardware specifications of the agents involved (Section 6).[6]

[6] Note that there is some work on learning from observations only (not actions) of a "teacher" [12]. However, this approach does not allow to integrate information such as the map in Example 3. Note that a difference to our method in Section 5 is that, e.g., an estimate of the complete transition probability is necessary, while our method only requires an idea of the "low-level" dynamics.

Generally, we aim at integrating information from many different sources simultaneously (e.g., many other self-driving cars in Example 1 and many different forms of information as described in Example 1 through Example 3. In particular, we aim at learning from databases that contain desirable as well as undesirable trajectories (e.g., avoid similar accidents as in Example 1).

Clearly, we do not present methodology that fully tackles the above shortcomings of LfD methods in this paper. Rather, we make first steps towards such methodology in Sections 5 and 6.

4.3 Multi-agent systems

In multi-agent systems, collections of agents acting in a shared environment are studied [15]. One important task is collaboration between agents [9]. A common approach is to model the collection of agents again as a single agent, by considering tuples of actions and observations as single actions and observations. Learning-based methods have been extensively studied [5].

While multi-agent systems approaches often allow to share and transfer information between agents, regarding the problem we formulated in Section 3 they have certain limitations: Similar as LfD, they usually do not integrate higher-level information sources (as the map in Example 3) or explicit hardware specifications of the agents (which we do in Section 6). Furthermore, if the mapping from some source agents sensory data to a target agents action is learned via modeling all agents as a single one, then it seems difficult to add agents to an environment, while our preliminary investigation in Section 6 in principle allows for adding agents more easily. Also note that the task of collaboration between agents is rather external to the problem we consider.

4.4 Transfer learning for agents

The problem we consider is related to transfer learning for agents. For instance, [17] consider an example where, in the well-known mountain car example, experience should be transferred although the motor of the car is changed. This comes close to transferring experience between self-driving cars as we suggest in Example 1. However, the scope of methods reviewed in [17] is on transferring observation-action recordings or things such as policies, value functions etc. using an appropriate mapping, while the goal we pursue is to also integrate information which is usually not expressible in these terms (e.g., the map or natural language description in Example 3, or the observation of another agent in Example 2). Furthermore, in this paper we aim at integrating *many heterogeneous sources* of information, while in transfer learning, even though several sources of information may be considered, they are usually homogeneous.

4.5 Further related areas

Other related directions include the following. Recently, the experimentation with intelligent agents in platforms based on computer games has become popular [8]. To our knowledge, the current work is the first one to use such platforms to study the problem of information integration, or related problems such as LfD (Section 4.2).

The general integration and transfer of data (not focused on intelligent agents) using causal models has been studied by [11, 3]. The idea of integrating higher-level information (again not for intelligent agents though) has been studied, e.g., by [18]. The relation between intelligent agents and causal models has been studied from a more philosophical perspective, e.g., by [19].

Another related areas is computerized knowledge representation [13]. Compared to general approaches to knowledge representation, our focus is on knowledge about the *physical* world.

Another relevant area is integration of human knowledge [7].

5 Investigation 1: integration of "non-subjective" information, evaluated in a simulated environment

In this section we aim to shed light on the following aspects of the problem formulated in Section 3:

- Generally, what experiments, in particular in simulated environments, can be performed to better understand the problem?
- How can experimentation (exploration) help an agent to translate "non-subjective" experience not recorded by itself into its own "coordinate system" and use it for (successful) decision making?
- How can partial information on the dynamics, such as a controller that is known to work locally, be merged with "higher-level" information such as hints on the path to some goal position?
- How can we quantify the efficiency gain from additional information sources?

The investigation is structured as follows: in Section 5.1 we describe the scenario, in particular the available heterogeneous information sources, in Section 5.2 we sketch an information-integrating agent for that scenario, then, in Section 5.3, we evaluate an adapted version of the agent in a simulated environment, and last, in Section 5.4 we discuss the advantages of our method over those that do not use additional information, and some further aspects.

5.1 Scenario

Task. An agent A starts in some unknown landscape and the task is to get to some visually recognizable goal position as quickly as possible.

Available heterogeneous information. We assume the following information sources to be available:

- the agent's own sensory input in the form of images y_t and position signal q_t (which can be seen an "interactive information" source since the agent can "query" this source via its actions),
- the controller ctl, which can be seen as a summary of the agent A's past subjective experience regarding the invariant local "physical laws" of a class of environments[7],
- a video trajectory $y^*_{0:L}$ that is a first-person recording of another agent with similar (but not necessarily identical) hardware that runs to the goal in the same environment.

Relation to the problem formulated in Section 3. On the one hand, this scenario can be seen as a (very) simplistic version of the scenario described in Example 1: A is a self-driving car that is supposed to get to some marked goal in an unknown environment, and the video $y^*_{0:L}$ comes from other cars that have a similar video-recording device but different hardware (engine etc.).

On the other hand, this scenario can be seen as a simplistic version of the scenario described in Example 3: the unknown landscape is some unknown city A arrived in, and instead of a description of the way to the destination in simple natural language, it gets a sequence of photos that describe the path it has to take.

[7] Specifically, we assume that if the distance between position q_1 and q_2 is small, then ctl successfully steers from q_1 to q_2. Alternatively, ctl could be a local model of the dynamics which induces such a controller.

Algorithm 1 Agent that integrates first-person video of other agent

input: Controller ctl, video $y^*_{1:L}$.
for i = 1, . . . , L **do**
  Use local controller ctl, optimization method opt and interaction with the environment to search locally around the current position for the next $q_i = \arg\min_q dist(y^*_i, \mathbb{E}(Y|Q = q))$.
  Use ctl to go to q_i.
end for

Algorithm 2 Proof-of-concept of Algorithm 1 for Malmo

input: Controller ctl, video $y^*_{1:L}$.
set $r_0 =$ current position, once the mission starts
for i = 1, . . . , L **do**
  use ctl, opt and teleportation to locally search around position r_{i-1} for the next $r_i = \arg\min_r \mathbb{E}(dist(y^*_i, Y)|Q = r)$
end for
restart the mission
set $i := 0$.
while $i < L$ **do**
  use ctl to steer to r_i
  if current position is close to r_i **then**
    set $i := i + 1$
  end if
end while

5.2 Method

First we sketch a general method, i.e., "software" for A, in Algorithm 1, assuming a (stochastic) optimization method opt and an image distance $dist$ as given (for concrete examples, see below). Note that in Algorithm 1 we denote by $\mathbb{E}(Y|Q = q)$ the mean image Y observed at position $Q = q$. The basic idea is that the agent uses local experimentation, based on prior knowledge of the local dynamics, to map the video $y^*_{1:L}$ into information (and eventually actions) that directly describes its own situation.

Although Algorithm 1 is in principle applicable to the experimental setup we consider in Section 5.3 below, we will evaluate Algorithm 2 instead, which is a simplified proof-of-concept implementation of it, making use of "teleportation", allowing the agent to directly jump to other positions without needing to navigate there. For Algorithm 2, as optimization method opt, we use simple grid search. (Note that instead one could use gradient descent or Bayesian optimization techniques.) Furthermore, we define the image distance $dist$ using Gaussian blur as follows: $dist(u, v) := \|N * (\bar{u} - \bar{v})\|$, where $\bar{u}$ is the normalization of u (i.e., subtraction by mean and division by standard deviation over the single pixels), N is a bivariate Gaussian with hand-tuned variance and $*$ is the convolution in both image dimensions.

5.3 Empirical evaluation in a simulated environment

To evaluate Algorithm 2, we consider three simple "Parkours" missions in the experimentation platform Malmo, described in Section 2. These missions consist of simple maps that have a special, visually recognizable, position which is defined as goal. A short description

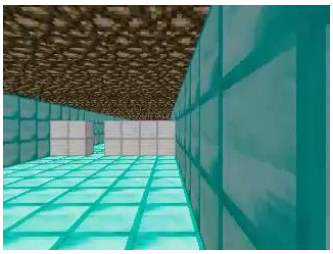

(a) Mission 1

(b) Mission 2

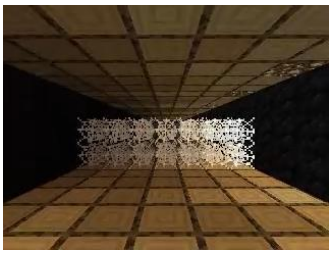

(c) Mission 3

Figure 1: First frames of the respective missions, for illustration purposes.

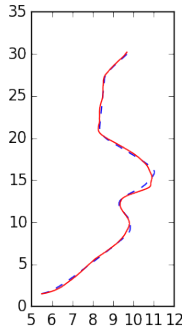

(a) Mission 1

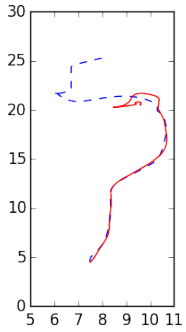

(b) Mission 2

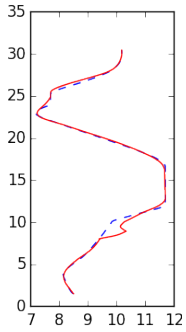

(c) Mission 3

Figure 2: The "ground truth" position trajectory of the demonstrator $q^*_{0:L}$ (blue dashed line), and the position trajectory of Algorithm 2 $\hat{q}_{0:K}$, (red solid line) from top view (x- and y-axis correspond to x- and y-coordinate in the map. While Algorithm 2 fails in Mission 2 due to the repetitive structure of some wall, it succeeds in Missions 2 and 3 in spite of its simplicity.

Table 1: For each mission (row) a short description (column 2), and the outcome of Algorithm 2 applied to it (columns three and four).

Mission	**Description and image**	U	$\hat{q}_{0:K}$ **versus** $q^*_{0:L}$
1	Figure 1a. Two passages have to be passed.	success	Figure 2a
2	Figure 1b. With (mortal) lava next to the path.	fail	Figure 2b
3	Figure 1c. With spider webs that occlude vision and lead to slow motion.	success	Figure 2c

of the three missions is given in column two of Table 1. We generally restrict the possible actions to $[-1, 1]^2$, where the first dimension is moving forth and back, and the second is strafing (moving sideways). The *task* is to get to the goal position within 15 seconds in these maps.

For each mission, we record one trajectory performed by a human demonstrator, which solves the task. More specifically, we record positions, which we denote by $q^*_{0:L}$, and observations (video frames), denoted by $y^*_{0:L}$.

We run Algorithm 2 with inputs $y^*_{0:L}$ and a simple proportional controller [2] (for ctl), where we tuned the proportional constant manually in previous experiments (but without providing $q^*_{0:L}$ or the actions the human demonstrator took). Let $\hat{q}_{0:K}$ denote the trajectory of positions that Algorithm 2 subsequently takes in the map. Furthermore, let $U \in \{\text{fail}, \text{success}\}$ denote whether the position tracking while-loop of Algorithm 2 (line 8 and following) gets to the goal within 15 seconds. (This is a significantly weaker evaluation metric than considering the runtime of the complete Algorithm 2 of course.)

The outcome of the experiment is given in columns three and four of Table 1 and Figure 2.

5.4 Discussion

5.4.1 *Discussion of the experiment*

Outcome. As shown by Table 1 and Figure 2, Algorithm 2 is successful for Missions 1 and 3. It fails in Mission 2 due to the repetitive structure of a wall that fills the complete image that is observed at some point during $y^*_{1:L}$. This wall makes the mapping from position to observation (video frame) (locally) non-injective which makes the algorithm fail. Note that this problem could quite easily be overcome by using prior assumptions on the smoothness of $q^*_{0:L}$ together with considering more than one minimum of $dist$ as the potential true position (using, e.g., Bayesian optimization) or by searching for position sequences longer than 1 that match $y^*_{1:L}$.

Limitation of the experiment. A clear limitation of the experiment is that the human demonstrator that produced $y^*_{1:L}$ used the same (simulated) "hardware" as Algorithm 2, while the overall goal of this paper is to integrate heterogeneous information. However, we hope that this experiment can form the basis for more sophisticated ones in the future.

5.4.2 *Theoretical analysis and insight*

Efficiency gain from harnessing $y^*_{1:L}$. Assuming there are at most N positions which the demonstrator can reach within one time step, Algorithm 1 takes only about $O(L \cdot N)$ steps to get to the goal. Note that this theoretical analysis is supported by the empirical evaluation: Algorithm 2's trajectories - visualized in Figure 2 - are roughly as long as $y^*_{1:L}$ (note that the visualization does not show the local search of length N).

This has to be contrasted with an agent that does not integrate the information $y^*_{1:L}$, and therefore, in the worst case, has to search *all* positions in the map, a number which is usually is much higher than $O(L \cdot N)$ (roughly $O(L^2)$).

Comparison to LfD. The task we study is closely related to LfD. However, note that usually in LfD, the target agent ("learner") has access to the demonstrator's actions, which is not necessary for our method. Furthermore, in our method, in some sense, the target agent can be seen as translating $y^*_{1:L}$ into its own "coordinate system" itself, while this mapping is usually hand-crafted in LfD.

Some insights during the development of the method. An interesting insight during the development of Algorithm 5.2 was that while the low-pass filter only led to minor improvements, what really helped was to visit each position say three times and then optimize the distance over the *averaged* images. Furthermore, it was was surprising how well the simple Euclidean distance we used worked.

Limitations. Note that the method is limited to environments similar to landscapes where some stochasticity and variation may be contained, but not too much. For instance, if the environment varies to strongly in the dependence on the time an agent spends in the environment, the proposed method most likely fails since the tracking of $y^*_{1:L}$ usually takes longer than the original performance of it.

6 Investigation 2: integrating sensory data, hardware specifications and causal relations

In this section we aim to shed light on the following aspects of the problem formulated in Section 3:

- How can information on the the hardware specifications of various agents be used for knowledge transfer between them?
- To what extent can *causal models* help, e.g., for integrating those hardware specifications (i.e., information on the "data producing mechanisms")?[8]
- How can information from the *"subjective perspective"* of an agent (i.e., on the relation between its sensory measurements and its actions) be merged with information from an "outside perspective" (i.e., that of an engineer which sees the hardware specifications of an agent).

The investigation is structured as follows: in Section 6.1 we describe the scenario, in particular the available heterogeneous information sources, in Section 6.2 we outline an information-integrating agent for that scenario, then, in Section 6.3, we give an intuitive toy example of scenario and method, and last, in Section 6.4 we discuss advantages and limitations of our approach.

Keep in mind the definition of causal models in Section 2. It needs to be mentioned that at certain points in this section we will allow ourselves some extent of imprecision (in particular in the treatment of the (causal) model M and how it is inferred), since we aim at going beyond what current rigorous modeling languages allow.

6.1 Scenario

Task. We consider a scenario where a collection C of autonomous agents, think of self-driving cars, operates in a shared environment. (For simplicity we assume that the number of agents is small compared to the size of the environment, such that they do not affect each other.) We assume that while some hardware components of the agents differ, others are *invariant* between them. We assume that for each car a task (e.g., to track some trajectory) is given and fixed.

Available heterogeneous information. Note that we could allow C to vary over time, e.g., to account for the fact that new cars get on the road every day, however, for the sake of a simple exposition, we leave it fixed here. We assume the following information sources to be available at time t:

[8] Another reason why causal reasoning could help is that in the end we are interested in the *causal effects* of an agent on its environment, and not just correlational information.

Algorithm 3 Integration and control algorithm for agent j

1: **input:** Time point t, description D, specifications $(spec^k)_{k\in C}$, experiences $(e^k_t)_{k\in C}$.
2: Initialize a causal model M by the causal diagram implied by the description D over the set of factors in $(spec^k)_{k\in C}$ and $(e^k_t)_{k\in C}$.
3: Update the "belief" over the mechanisms in M using all values of variables contained in $(spec^k)_{k\in C}$ and the experiences $(e^k_t)_{k\in C}$ (possibly based on additional priors).
4: From the updated M, calculate M^j, the implication of M for agent j.
5: Find action $u(t)$ that is optimal w.r.t. the given task, under M^j.
6: **output:** $u(t)$

- specifications $spec^k$ of the hardware of each agent k
- past experiences (i.e., actions and observations) e^k_t of all agents $k \in C$, consisting of observations $y^k(t)$ and control outputs $u^k(t)$, i.e., $e^k_t = (u^k(0), y^k(0), \ldots, u^k(t), y^k(t))$.
- a description D (e.g., a physical or causal model or collection of such models) consisting of (1) a set of *independence statements*[9] and (2) a set of dependence statements, possibly including specific information on the shape of the dependence, w.r.t. the factors contained in the specifications $spec^k$ and the experiences e^k, for $k \in C$. Potentially, the various independence and dependence information pieces could come from different sources, say targeted experiments as well as general prior knowledge. We assume the dependence structure (including the precise shapes of the dependences) to be time-invariant.

Relation to general problem. This scenario captures certain aspects of Example 3 in that some higher-level information in the form of the description D is available. While here, as a first step, we only consider mathematical models, in the future one could also imagine to include informal but well-structured models and descriptions (in simple natural language), possibly translating them into formal models as an intermediate step (using, e.g., machine learning). Furthermore, the scenario captures important aspects of Example 1, since we consider the integration of information from certain source self-driving cars for the decision making of a given target car.

6.2 Sketch of a method

We sketch a method for the described scenario in Algorithm 3. It first derives a "global" causal diagram - applying to all agents - from the potentially heterogeneous description D (line 2). Then (line 3) the causal conditionals of M, which are not determined by D, are inferred from the given hardware specifications as well as the experiences gathered by all agents up to time t. Last, based on the hardware specifications of agent j, the implications of M for agent j are calculated (line 4) and the optimal action under these implications is performed (line 5).

6.3 A toy example

Let us illustrate how the method proposed in Section 6.2 works in a concrete toy scenario. The core intuition is that while some details

[9] Clearly, independence assertions are central to integration of information: only based on statements of the form "Y is more or less independent of all factors that potentially will be included, except for this and this small set" it seems possible to rigorously (automatically) reason about integration.

of the dynamics of self-driving cars may vary between different cars, they can still share information on say the road conditions (friction, drag, etc.) at certain positions y, or the like.

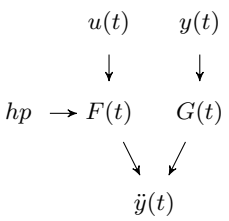

Figure 3: Sketch of the causal diagram H. The power hp influencing only $F(t)$ implies that knowledge on the mechanism f_G for $G(t)$ can be transferred between two cars even if they differ in hp.

Specific scenario. We consider two simplified self-driving cars, i.e., $C = \{1, 2\}$, where we assume the observation $y^k(t)$ (or $y(t)$ if we refer to a model for both cars) to be the car's position. We consider the following concrete instances of the information sources listed in Section 6.1. (Note that, for the sake of simplicity, we assume all mechanisms to be deterministic, such that we can model them by functions $f_{...}$ instead of conditional distributions $p(\dots|\text{pa}_{...})$, although the latter would be more general, see Section 2.)

- The specification for car k is given by its power (e.g., measured in horse powers), i.e., $spec^k = hp^k$.
- The experience of car k consist of all past position-action pairs of both cars, i.e., $e_t^k = (u^k(0), y^k(0), \dots, u^k(t), y^k(t))$.
- The description D consists of three elements:
  - an engineer contributes the equation $F(t) = f_F(u(t), hp)$ for the engine of the cars[10], where $F(t)$ is the force produced by the engine (incl. gears), and f_F a known function,
  - a physicist contributes the equation $\ddot{y}(t) = \frac{1}{m}(F(t)+G(t))$ for the acceleration $\ddot{y}(t)$ of the cars, where $G(t)$ are other forces that affect the cars, such as friction and drag, and m is the known mass (which we assume to be the same for both cars),
  - another physicist produces a set of additional independence assertions, such that altogether the description D implies the causal diagram H depicted in Figure 3.[11]

Implementation of our method. We suggest the following concrete implementation of the crucial part of our method, i.e., line 3 and 4 in Algorithm 3, based on the concrete instances of information available in this toy scenario. Keep in mind that f_G, the function that maps a position y to the corresponding force G at that position and thus models the generating mechanism for $G(t)$, is the only unknown part of M after initializing it by D.

- Line 3: Use the experience $(e_t^k)_{k\in C}$, to infer the function f_G on all positions y visited by either of the cars, based on the equation

$$m\ddot{y} - f_F(u, hp^k) = f_G(y)$$

for all $k \in C$, and the fact that the l.h.s. of this equation as well as y are known for all positions (and accelerations) visited by either of the cars.
- Line 4: calculate "$p(\ddot{y}|do(u), y, hp^j)$", i.e., the effect of control action u of agent j at position y, *for all positions that were visited before by either of the cars*.

6.4 Discussion

The toy example in Section 6.3 shows how in principle the integration of heterogeneous information could help some "target" self-driving car for better decision making in situations not visited by it but by different "source" self-driving cars. It is important to note that all listed information sources were necessary for this: the hardware specifications are necessary to understand $F(t)$, the experience is necessary to infer f_G, and the independence knowledge (hp not affecting $G(t)$) is necessary to transfer the knowledge about the force $G(t)$ on various positions y between the cars. Note that the above scenario cannot be tackled by standard RL approaches since we transfer knowledge between agents of different hardware. Furthermore, methods like LfD or transfer learning (see Sections 4.2 and 4.4) usually do not automatically harness information on hardware specifications of agents.

Based on our preliminary investigation above, it seems that causal models are helpful in that they provide a language in which one can express relevant assumptions and reason about them. However, from a practical perspective, it is not clear if the necessary calculations could not be genuinely done e.g. in classical probabilistic models.

An important question is whether the method sketched in Section 6.2 can be generalized to dependence statements in more natural - but still well-structured - language than equations and causal diagrams.

7 Outlook: future directions and open questions

Here we sketch a potential agenda for future investigations and pose interesting open questions.

7.1 Potential future directions

"Universal representation of physical world". An interesting subject-matter of future research would be a "universal representation" of the physical world - a rich representation to which each information source could be translated, and from which each agent could derive the implications for its specific sensor and actuator configurations. Such a representation would be more efficient than handcrafting mappings for each (new) pair of source and agent individually (as is usually done in e.g. LfD, see Section 4.2), reducing the number of necessary mappings from n^2 to n, where n is the number of agents. One starting point would be representations that are already used to integrate laser or radar scanner data on the one hand with (stereo) video camera data on the other hand in self-driving cars [6]. Another starting point would be the global positioning system (GPS) which is a successful universal representation of position with clear, hardware-independent semantics.

Investigation and classification of the "integration mapping". Another important concept for integration of heterogeneous information could be the *mapping* that transforms a collection of pieces of well-structured heterogeneous information into a model of the current situation or even directly into action recommendations. The study of such a mapping could build on the investigation of related

[10] In particular, the mechanism specification implies non-influence by all other relevant factors.

[11] A better way to describe the physical forces causally might be to replace $\ddot{y}(t)$ in the equations and in D by expressions based on $v(t + \Delta t)$, $v(t)$, and Δt, where $v(t)$ denotes the velocity.

mappings in LfD and transfer learning for agents (see Sections 4.2 and 4.4). Furthermore, parts of such a mapping could be learned (which would be related to machine-learning-based multi-agent systems, see Section 4.3. Generally, it would be interesting to examine the basic conditions under which the integration of heterogeneous information can be beneficial. Note that one way to *classify and order* various sources of information (and the mappings that are necessary to integrate them) would be from "closest to the agent A", i.e., it's own past observations and actions to "most distant", e.g., agent-independent descriptions of the world in simple natural language, as exemplified in Section 3.

Further experiments. Further experiments, with a gradually increasing difficulty (e.g. along the ordering proposed in the previous paragraph), could be performed, e.g. using the platform Malmo (Section 2) to gain a better understanding of integration of information:

1. Agents can observe other agents from a third-person perspective, enabling Example 2 in Section 3.
2. Higher level observations can be provided in the form of natural language (typed chat or external information in the form of natural language), or through artifacts such as maps, sign-posts, symbolic clues, etc.

7.2 Open questions

It would also be interesting to investigate how the following questions could be answered:

- One of the main question which guided our investigation in Section 6 can be cast as follows. While the information relevant to an agent is usually in the form of effects of its actions in certain situations, a lot of knowledge is formulated in non-causal form: for instance street maps at various granularities for self-driving cars. How are these two forms of information related? Is there a standard way to translate between them? Stated differently, how can various forms of information be translated into a model of the *dynamics* of the agent in the world.
- Where is the boundary between additional heterogeneous information and prior knowledge?
- How can the need for information integration be balanced with *privacy* restrictions? For instance, one may imagine cases where the mapping from a source agent's experience to a target agent's action is rather simple in principle, but information collected by the source agent cannot or should not be transmitted to the other, at least not in full.
- How can big databases of information be filtered for useful information, i.e., the information which is correct and relevant for the current environment and task?
- To what extent is the problem of information transfer between two different agents in the "same" environment just a special case of transfer between different environments (by considering the hardware of an agent as part of the environment)?
- How can we reason without having a *"global" model* such as Figure 3? What about interfaces to build global from *"local" models*, describing only say the engine?
- Generally, what are potential theoretical limitations of automated information integration, e.g. in terms of computability?

8 Conclusions

In this paper, we considered the problem of designing agents that autonomously integrate available heterogeneous information about their environment. We investigated how experimentation in simulated environments on the one hand, and causal models on the other, can help to address it. A next step would be to perform more sophisticated experiments, ideally guided by specific problems e.g. from the area of self-driving cars.

9 Acknowledgments

The authors thank Mathew Monfort, Nicole Beckage, Roberto Calandra, Matthew Johnson, Tim Hutton, David Bignell, Daniel Tarlow, Chris Bishop and Andrew Blake for helpful discussions, and the anonymous reviewers for useful hints.

REFERENCES

[1] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning, 'A survey of robot learning from demonstration', *Robotics and Autonomous Systems*, **57**(5), 469 – 483, (2009).
[2] Karl Johan Aström and Richard M Murray, *Feedback systems: an introduction for scientists and engineers*, Princeton university press, 2010.
[3] Elias Bareinboim and Judea Pearl, 'Causal inference from big data: Theoretical foundations and the data-fusion problem', Technical report, CALIFORNIA UNIV LOS ANGELES DEPT OF COMPUTER SCIENCE, (2015).
[4] David Bignell, Katja Hofmann, Tim Hutton, and Matthew Johnson, 'The Malmo platform for artificial intelligence experimentation', in *IJCAI (to appear)*, (2016).
[5] Lucian Buşoniu, Robert Babuška, and Bart De Schutter, 'Multi-agent reinforcement learning: An overview', in *Innovations in Multi-Agent Systems and Applications-1*, 183–221, Springer, (2010).
[6] Andreas Geiger, Julius Ziegler, and Christoph Stiller, 'Stereoscan: Dense 3d reconstruction in real-time', in *Intelligent Vehicles Symposium (IV), 2011 IEEE*, pp. 963–968. IEEE, (2011).
[7] Andreas Holzinger, 'Interactive machine learning for health informatics: when do we need the human-in-the-loop?', *Brain Informatics*, **3**(2), 119–131, (2016).
[8] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al., 'Human-level control through deep reinforcement learning', *Nature*, **518**(7540), 529–533, (2015).
[9] Reza Olfati-Saber, Alex Fax, and Richard M Murray, 'Consensus and cooperation in networked multi-agent systems', *Proceedings of the IEEE*, **95**(1), 215–233, (2007).
[10] J. Pearl, *Causality*, Cambridge University Press, 2000.
[11] Judea Pearl and Elias Bareinboim, 'Transportability of causal and statistical relations: A formal approach', in *In Proceedings of the Twenty-Fifth National Conference on Artificial Intelligence. AAAI Press, Menlo Park, CA*, pp. 247–254, (2011).
[12] Bob Price and Craig Boutilier, 'Accelerating reinforcement learning through implicit imitation', *J. Artif. Int. Res.*, **19**(1), 569–629, (December 2003).
[13] John F Sowa, 'Knowledge representation: logical, philosophical, and computational foundations', (1999).
[14] P. Spirtes, C. Glymour, and R. Scheines, *Causation, prediction, and search*, MIT, Cambridge, MA, 2nd edn., 2000.
[15] Peter Stone and Manuela Veloso, 'Multiagent systems: A survey from a machine learning perspective', *Autonomous Robots*, **8**(3), 345–383, (2000).
[16] Richard S Sutton and Andrew G Barto, *Reinforcement learning: An introduction*, MIT press, 1998.
[17] Matthew E Taylor and Peter Stone, 'Transfer learning for reinforcement learning domains: A survey', *The Journal of Machine Learning Research*, **10**, 1633–1685, (2009).
[18] Vladimir Vapnik and Akshay Vashist, 'A new learning paradigm: Learning using privileged information', *Neural Networks*, **22**(5), 544–557, (2009).
[19] James Woodward, *Making things happen: A theory of causal explanation*, Oxford University Press, 2005.